

UNITED STATES PATENT APPLICATION FOR
ADDITIVE MODEL FOR EFFICIENT REPRESENTATION OF
DIGITAL DOCUMENTS

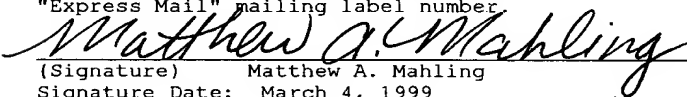
Inventors:

Doron Kletter
Donald James Curry

CERTIFICATE OF MAILING BY "EXPRESS MAIL"
UNDER 37 C.F.R. §1.10

"Express Mail" mailing label number: TB597446168US
Date of Mailing: March 4, 1999

I hereby certify that this correspondence is being deposited with the United States Postal Service, utilizing the "Express Mail Post Office to Addressee" service addressed to **Box PATENT APPLICATIONS, Assistant Commissioner for Patents, Washington, DC 20231** and mailed on the above Date of Mailing with the above "Express Mail" mailing label number.


(Signature) Matthew A. Mahling
Signature Date: March 4, 1999

ADDITIVE MODEL FOR EFFICIENT REPRESENTATION OF
DIGITAL DOCUMENTS

Inventors:

Doron Kletter
Donald James Curry

5

BACKGROUND OF THE INVENTION

Field of Invention

10 This invention relates to the representation of
digital images, particularly the representation of
those digital images in compressed format. The
invention is also related to the representation of
digital images in multiple planes of like, similar, or
related data traits or values. The invention is more
particularly related to the optimization of a plane of
15 a digital image for a selected compression methodology
and for storing correction data for the optimized plane
in an alternate plane of the image. The invention is
also related to the reconstruction of a representation
of a digital image by selecting data from either an
20 optimized plane of a representation of the digital
image and an additive plane comprising additive data
from each of the optimized plane and a correction
plane.

Discussion of Background

The efficient transmission, exchange, and storage of image data is a key issue in distributed systems. The main difficulty is to find effective ways of significantly reducing the large amount of information that needs to be transferred and/or processed, while still being able to reconstruct the image with good fidelity. The conventional approach to deal with this problem is to reduce the amount of data by applying compression. Several compression technique (such as LZW, JPEG, JBIG, and Wavelets) have been developed over the years to effectively handle the various types of data often found in digital documents. For example, binary compression schemes such as LZW are very effective in dealing with text and line-art information, while JPEG has been successfully developed to encode natural photographs. However, these algorithms have been designed to perform well on a particular class (be it textual or pictorial) of images - and no single algorithm can possibly handle all types of image classes reasonably well.

One approach to achieve high compression ratio is to divide the underlying document by image class, and apply the most appropriate compression technique on each class separately. Text and line-art information would be compressed using an approach that focuses on

maintaining the detailed structure (edge information), while pictures and color sweeps would be compressed using an approach that opts for preserving the color depth and smoothness.

5 The Document Image Representation (DIR) is one way of describing digital documents based on this concept. According to the DIR draft specification (Rev. 1.1 dated August 26, 1996), the current DIR model is represented using three different (logical) planes, as shown in Figure 1: An upper plane Up, also referred to as the foreground plane, typically containing data that puts high requirements on detail such as text and line-art; A purely binary selector plane Sp, or mask plane, that provides switching information between the upper and lower planes; And a lower plane Lp, also referred to as the background plane, that emphasizes color gradation, such as images and sweeps. Note that the contents of the planes in Figure 1 is shown just for the purpose of illustration. One could, for example, represent the same image in Fig. 1 (colored text surrounding a picture) by having binary (black) text in the selector plane, and filling the upper plane with the color of the text everywhere. Other variations are also possible.

25 The embedded imaging model within DIR will be referred to as the Selective Model, since content of

the upper plane is used to overwrite the content of the lower plane, on a pixel by pixel basis, whenever the selector plane is true. In other words, the upper plane data is "poured" through the binary selector plane, to
5 replace the lower plane data below. Note that at the end, each reconstructed pixel carries information from either the upper or lower planes. That is, the information present in the other plane (the one not selected by the mask) is considered irrelevant and is
10 ignored.

Adobe Systems, Inc., is currently considering to incorporate a new Masked Image Operator as an extension to Postscript Level III and PDF that could enable DIR reconstruction from a logical multi-planar
15 representation. The new operator can "mask" out the background, leaving a foreground object visible, without having to define a clipping path. The resulting image can then be "placed over" (e.g., overwrite) another background. Hence the Masked Image Operator
20 shares the same property as the embedded imaging model proposed in the DIR specification: they both share the notion of "pouring" the foreground through the mask, and overwriting the data below.

Typically, in each of the DIR and Adobe methods,
25 the upper and lower planes are usually compressed using very different compression schemes. The DIR

specification recommends, as an example, use of the following algorithms (DIR spec, pp.8): Fiala-Green (LZW related) for text and line-art (Upper plane), CCITT group 3/4 for the binary mask (Selector plane), and
5 JPEG for color images and sweeps (Lower plane). Hence with the Selective Model, one is forced to pick, per pixel, one compression scheme at all times and is prevented from using a clever combination of both. Even if one could potentially decide on how to separate
10 the information between planes, such decision may inevitably lead to loss of image quality, or alternatively, to comparable reduction in compression. For example, true line-art will not retain a good edge definition when passed through JPEG, unless the
15 compression requirements were greatly relaxed. Similarly, LZW will fail to produce good compression results when applied to noisy photographs.

SUMMARY OF THE INVENTION

The present inventors have realized that a
20 decision to enforce one of the planes at all times limits the flexibility of a representation model.

Accordingly, it is an object of the present invention to provide an image compressor and a method of image compression/representation that varies an

amount of detail contained in any of plural image planes representing an image.

It is another object of the present invention to optimize at least one of the plural image planes for
5 compression using a selected compression methodology by varying an amount of image detail maintained in each of the image planes.

It is yet another object of the present invention to determine an amount of image detail in each of the
10 image planes based on at least one of a super-resolution, a fine edge positioning pixels of a selected image plane, and an average color of the selected image plane.

It is still yet another object of the present
15 invention to provide an image decompressor and a method of image decompression that reconstructs an image from plural image planes, including an upper plane and an additive plane that is constructed from image information contained in both the upper plane and an
20 alternate plane the plural image planes.

These and other objects are accomplished by a method of decompressing an image, including the steps of decompressing an upper plane of an image, decompressing a lower plane of an image, adding the
25 decompressed upper plane and the decompressed lower plane to produce an additive image, and selecting

pixels of a final decompressed image from one of the decompressed lower plane and the additive image. The step of selecting pixels may include the step of varying an amount of the additive image produced by the upper plane based on a predetermined factor.

In addition, the above objects may be accomplished via a compression apparatus, including means for compressing an upper plane of an image and a lower plane of an image, wherein an amount of detail of the image contained in pixels of the compressed upper plane increases with respect to a predetermined factor of corresponding pixels in the image. In alternate embodiments, the predetermined factor may be one of a degree of color in the corresponding pixels matching a background color of the image and an amount of fine edge positions in the corresponding pixels.

The Document Image Representation (DIR) specification provides a way of describing digital documents by separating textual (line-art) information from continuous-tone pictures and applying different compression on each component, to take advantage of its specific attributes. However, the imbedded imaging model (a Selective Model) as implied by the DIR is limiting: the information must always be classified and strictly carried in one of two channels, but never in both.

The present invention improves and further enhances the existing DIR representation to eliminate the above discussed limitations. Specifically it offers to replace the Selective Imaging Model with an Additive
5 Imaging Model, thereby allowing the flexibility to combine information from both channels to make the best use of available compression resources. In addition, it offers to optionally extend the selector plane beyond binary. These modifications define a new and more
10 flexible method of representing digital documents, that may lead to substantial improvements both in quality and compression results when compared with the existing method. The improvements are demonstrated in a case study example discussed hereinbelow.

15 BRIEF DESCRIPTION OF THE DRAWINGS

The file of this patent contains at least one drawing executed in color. Copies of this patent with color drawings will be provided by the Patent and Trademark Office upon request and payment of the
20 necessary fee.

A more complete appreciation of the invention and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following detailed description when

considered in connection with the accompanying drawings, wherein:

Fig. 1 is a diagram illustrating the upper plane, selector plane, and lower plane of a document image representation model;

Fig. 2 is a diagram illustrating the selection process between an upper plane and a lower plane of a selective model reconstruction;

Fig. 3a is a diagram illustrating the selection between an additive plane and a lower plane of the additive model reconstruction according to the present invention;

Fig. 3b is a diagram illustrating the selection between an additive plane and a lower plane, each modified according to a detect signal in an alternate embodiment of the additive model reconstruction according to the present invention;

Fig. 4 illustrates a tile pattern utilized in illustrating the processes of additive model reconstruction according to the present invention;

Fig. 5a illustrates the three planes in a selective model reconstruction where all image information is contained in a lower plane only;

Fig. 5b illustrates the three planes of a selective model reconstruction having a single color

(red) maintained in an upper plane and all remaining information maintained in the lower plane;

Fig. 5c illustrates the three planes of a selective model reconstruction having a single color (green) maintained in the upper plane and all additional information being maintained in the lower plane;

Fig. 5d illustrates three planes of the selective model when the upper plane maintains a noisy background color and a lower plane maintains foreground color information;

Fig. 5e illustrates the three planes of the selective model representation where the upper plane maintains the foreground colors and the lower plane maintains the noisy background color;

Fig. 5f illustrates the three planes of a selective model representation when the upper plane maintains all information contained in the image;

Fig. 6a illustrates the first step of separation of planes according to the additive model of the present invention;

Fig. 6b illustrates the second step of color replacement in an upper plane according to the additive model of the present invention;

Fig. 6c illustrates the final step of embedding errors from upper plane color replacement into a lower

plane according the additive model of the present invention;

Fig. 7 illustrates a luminance of the selective model representation according to alternative 2b on
5 Fig. 2;

Fig. 8 illustrates an average luminance according to the additive model of the present invention;

Fig. 9 illustrates a source image;

Fig. 10a illustrates a selective model upper plane
10 for the source image of Fig. 9;

Fig. 10b illustrates a selector plane for the selective model in accordance with the upper plane of Fig. 10a and the source image of Fig. 9;

Fig. 10c illustrates a selective model lower plane
15 for the source image of Fig. 9;

Fig. 11a illustrates an additive model upper plane for the source image of Fig. 9 according to the present invention;

Fig. 11b illustrates a selector plane for the
20 additive model in accordance with the upper plane of Fig. 11a and the source image of Fig. 9 according to the present invention; and

Fig. 11c illustrates an additive model lower plane
25 for the source image of Fig. 9 according to the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring again to the drawings, wherein like reference numerals designate identical or corresponding parts, and more particularly to Fig. 2 thereof, there is illustrated the existing DIR model based on a Selective Imaging Model. In this model, the output image 20 is reconstructed from the DIR representation (the three planes described above, comprising of an upper, lower, and binary selector plane) using a simple multiplex operation. The binary selector value 22, derived from the selector plane, is used to choose data from either the upper plane 10 or lower plane 12, on a pixel by pixel basis. Note that any information that might be present in the other plane (the one not selected by the mask) is ignored and treated as irrelevant.

As already explained, the problem with this model is that a decision must be made for each and every pixel: to select one of the upper or lower planes. The DIR spec does not specify a recommendation as to how to make that decision, so in theory full freedom exists. However, the DIR specification does make a recommendation to use certain standard block based compression techniques for some of the planes (e.g., the use of JPEG for compressing the lower plane). This indirectly implies that the pixels in a block are tied

together via the compression, regardless of how they are classified.

A problem occurs whenever not all the pixels within a block are classified to be of the same type. Perhaps the most obvious example that comes to mind is the case where there is some (noisy) text on top of an image. A block of pixels situated on the text boundary will occasionally contain a portion of the text as well as a portion the image. In a perfect world, the text portion could be "lifted" off the image to be encoded in the upper plane, and the remainder is left to be JPEG'ed in the lower plane. However, because pixels within a block are inter-dependent by the compression, the lifted text pieces may end up affecting the quality or compression ratio of the non-textual areas nearby.

Furthermore, the current DIR specification recommends the use of a variant of LZW for compressing the upper plane. This technique favors binary situations with few colors in order to achieve high compression ratios. If several pixels within a block are classified as line-art, and all have different colors (possibly due to noise variations), then the resulting compression performance may drop substantially (as will be demonstrated by the example below).

In this case, one is caught in a fundamental conflict resulting from the inflexibility of the Selective Model: either limit the number of pixels classified as line-art (to keep compression in check),
5 or force them all to have the same color(s). Limiting the number of line-art pixels may result in quality degradation, for the other line-art pixels (the ones not classified as such due to the limit) will be sent through JPEG and may now lose the edge detail.
10 Similarly, if the color(s) are forced, an immediate representation error is made; This error will be carried forward to the output, since a pixel in the upper plane overrides the lower, with no ability to correct for it at a later stage. Hence the existing DIR
15 Model lacks the ability to compensate for desirable adjustments (such as due to compression-related constraints) in the way pixels are classified.

As an alternative, consider the new Additive Imaging Model proposed in this invention as illustrated
20 in Fig. 3a. This model has the advantage in reconstruction, whereby information from an upper plane 32 is first (prior to making the selection) added to the content of a lower plane 36. Note that a significant difference from the previous Selective
25 Model is the introduction of adder 38. The adder 38 combines data from the two channels (Upper plane 32 and

lower plane 36), thereby providing a mechanism to correct for representation errors resulting from the classification adjustments discussed above.

The adder may be implemented in alternative ways.

5 For example, the adder may be utilized to combine pixels from the upper and lower planes at the time of selection. Alternatively, the adder may be utilized to combine the upper and lower planes to produce an additive image, and then the selector chooses pixels
10 from either one of the upper plane (in this embodiment) and the additive image. As will be appreciated by one skilled in the art, other configurations to procure the same end result are also possible (selections based on blocks or clusters, as one example).

15 The Additive Imaging Model can significantly improve the results. For example, given the type of upper plane compression used, the foreground color may be forced to be composed of a limited number of colors (a decision likely to substantially improve the
20 compression), or even to "not-so-close colors", yet compensate for this via the background channel (lower plane in this example). We do so by calculating the individual pixel errors resulting from such a decision, and then subtracting this error from the background
25 value. Furthermore, we can also individually (on a pixel by pixel basis) control how much of the modeling

error we wish to carry in each of the channels. This enables compression and quality to be optimized together. It is a huge benefit for the cost of adding a single adder (which may be implemented in either software or hardware).

The Additive Model can also be naturally extended to provide greater resolution for edge detail. Current practice seems to indicate that having a selector plane at higher resolutions is desirable whenever the input resolution is below 600 dpi. For this range, typical Resolution Enhancement techniques (Ret) have been successfully employed to anti-alias (smooth) text and line-art to improve its appearance.

Alternatively, a non-binary selector plane may be utilized. With more than one bit, the selector plane naturally extends to become an alpha channel for mixing (or blending) information between the upper and lower planes. Of course, one would want to limit the number of bits allocated to the selector plane for better compression, but in some situations additional selector plane bits may be advantageous.

The alternative embodiment (extension) is illustrated in Fig. 3b. A detector 38 interprets a value and read from a multi-bit selector plane which is utilized to alter (see amp/mixers 39a and 39b) values

read from each of an upper plane 33 and lower plane 37 prior to producing an additive plane (via adder 40).

With this extension, the extra bits in the selector plane are used as weights to vary the mixing proportions of the two channels. One useful application of such non-binary selector could be super-resolution, or fine edge positioning. The closer the edge is, the more weight is given to the foreground as opposed to the background. This technique can eliminate the need for having REt (as a post-processing step) to enhance text and line-art appearance.

To illustrate the advantages of the new approach, consider the example pattern shown in Figure 4. Suppose the input image is composed of a page that is fully covered with the regular pattern shown (the motivation for choosing this example will become apparent from the description below). The pattern is based on a basic four-quadrant cell whose building element is encapsulated within the center dark frame 40 in Figure 4. This cell is then regularly tiled over the page to create the complete pattern.

This example, as illustrated in Fig. 9, demonstrates that in some cases the new technique can produce significantly better results overall, in terms of the flexibility of the model, the fidelity of representation, and the bottom line compression ratio.

True, this example is constructed here merely to make a point, not necessarily to establish a global conclusion. However, as will be demonstrated in this section, the differences can be substantial - enough to
5 justify enabling the new technique.

Assume, without any loss of generality, that the tile size is 8 x 8 pixels at the given input resolution. This will establish a relationship between the chosen cell size and the standard JPEG block size
10 for the purpose of evaluating compression performance (otherwise the cell size can be considered to be related to the local area of interest with respect to the given viewing conditions and input resolution). It then follows that each quadrant is composed of a total
15 of (4 x 4) = 16 pixels; Two of the quadrants, the red and the green, are collectively referred to as "foreground colors" 42 (unrelated to the plane- names below in any way). Assume further that little noise is present in the foreground (as would be the case
20 whenever images are scanned), and therefore all the pixels in each of the foreground quadrants have very close color values, but not necessarily the same. Similarly, the remaining two quadrants (shown gray) are referred to as "background color" 44 and contain
25 portions of a different contone (continuous tone) image (likely to be noisy too).

The main idea here is to have a relatively "clean" two-color foreground texture on top of some noisy background image, similar to what one would get when scanning colored text or line-art on top of an image.

5 This example is extreme in the sense that every image cell contains a mixture of dual-color foreground and background, a situation not likely to occur very often. In practice, the number of such mixed cells will vary widely, depending on page complexity and segmentation
10 details. Nonetheless, useful insight can be learned from this example.

With a 4 x 4 pixel size, the quadrants are generally large enough to be quite noticeable (at a typical practical resolution of 300 to 600 dpi). Note
15 that there are three substantially different color areas altogether in Figure 4 (the example will work as long as the number of distinct foreground colors is greater than one). The pattern of interest was generated using Adobe Photoshop for the purpose of
20 evaluating the compression results. The method of generating the pattern was by filling an 8 x 8 pixel area with the two solid red/green foreground colors, leaving the two background quadrants as a transparent layer. This area was then tiled to fill an image of the
25 same size as the input image, and independent random Gaussian noise was added everywhere (for all pixels).

Finally, the pattern image was applied (using a mask operation) to a given background image (a scanned image of "Lenna", cf. Figure 9) of equal size such that the background data was propagated through the "holes" in the mask.

Figures 5a - 5f demonstrate some of the possible ways for representing this tile image according to a DIR representation compatible with the current specification (based on the existing Selective Model).

Consider case 5b), for example; Here, the red quadrant information is classified to be in the upper plane, while the three other quadrants are left in the lower plane. The yellow squares indicate areas of irrelevant information in the content of both the upper and lower planes. Since it does not matter what is in there, the best utilization of this space is to fill it in with values that make the corresponding compression as tight as possible (e.g., for case 5b above, one might as well fill the whole upper plane with red everywhere, since this will undoubtedly lead to the highest compression ratio).

The examples shown in Figures 5a and 5f do not make sense; The selector value is set to either true or false everywhere, forcing the selected (single) plane to carry all the image information; And if the image happens to be complex (as is the case for the tile

pattern of interest), one will not be able to effectively compress it, artifact free.

The case illustrated in Figure 5b is very similar to 5c (a simple color in the upper plane). However, in terms of compression, they may still differ: even though they both share a constant value upper plane (good compression), the lower plane (JPEG) may still compress differently. Obviously, the preferred choice is to pick the one color "red" or "green") that is closest to the "gray" as part of the background. But if both are significantly different from the "gray" such that a strong edge exist on the borders, either case will lead to significant [JPEG] "ringing" artifacts - unless the compression requirements are greatly relaxed. Cases 5b and 5c, however, also share an undesirable quality downside: they both are non-symmetric in the sense that they treat the red and green colors unevenly. The color that is in the upper plane [LZW] will normally maintain uniform density and sharp edges, while the other [JPEG] will vary in density and edges will "ring."

The case illustrated in Fig. 5e may lead to better overall compression than 5a and 5b with respect to both the upper and lower planes combined. None-the-less, it has a more complex selector structure (checkerboard pattern) that might not be as effective to compress

(with Lossless fax Group 4 CCITT). If both the red and green are sufficiently dark to be true foreground, the preferred representation might well be the one in case 5e, for the selector plane in this case truly
5 represents the binary mask.

Finally, case 5d is like 5e, except the role of the foreground and background planes are reversed. Again, for a given choice of compression pair (i.e., LZW and JPEG) tied to planes at hand, one of these
10 cases may perform significantly worse. Similar "role-reversal" situations also exist for cases 5b and 5c, but are not shown.

The main purpose of this discussion regarding Figures 5a - 5f is to demonstrate that the tile pattern
15 in Figure 4 is complex enough such that one cannot find a simple "artifact free" representation using the Selective Model (without giving up compression). No matter how the data is split, you always have at least one of the upper or lower planes containing alternating
20 checker patterns that will not compress well (resulting in short run lengths in LZW, or edge ringing in JPEG). In fact, the above will hold true for any tri-tone case where more than one "foreground" color exists.

We will now demonstrate how the new Additive
25 Imaging model hereby proposed can significantly improve the results. Consider the preferred case illustrated

in Figure 5e. Here, the lower plane poses no problem,
but the upper plane is checkered with at least two
colors (red/green). To make it compress better, suppose
we attempt to replace the two colors with just a single
5 color. We can pick one of the two, or we can use their
average. This will certainly improve the compression of
the upper plane, but will also introduce a substantial
representation error (a single color instead of the
patterned two). With the Selective Model, the error
10 will be carried through, leading to unacceptable
results (the only option is to use the lower plane with
low compression). With the Additive model, however,
this error can be fed back to adjust the lower plane.
At reconstruction time, the information from the two
15 planes will be combined (added) together to regenerate
the source image.

To illustrate the technique, consider the three
steps shown in Figures 6a - 6c. In the first step,
Figure 6a the input image is separated into the three
20 planes (conceptually similar to Figure 5e). One
difference being that the lower plane information is
always used (due to the additivity) and therefore must
be set to some meaningful value (typically the average
background color).

25 Next, in step 2, Figure 6b, the red and green are
replaced by a single color (the average foreground)

that best represents both colors (denoted as dark gray color in Figure 6b and 6c). The difference between this average and each of the two colors is then added to the lower plane in step 3 as seen in Figure 6c. The color coding scheme illustrated is that the inverse of red is cyan, and the inverse of green is magenta. Note, however, that if both "red" and "green" colors have been identified as foreground, then they are generally of significant contrast and would typically be closer to each other (in luminance) than to the background. The three planes are then compressed as before.

A one-dimensional, luminance only, illustration of the above procedure is illustrated in Figure 7. Imagine looking at a luminance projection of Figure 6a. There is a (slowly varying) background level on the lower plane, and alternating red/green checkers on the upper plane. The content of selector plane is not shown. Under the new technique, both the red and green luminance are replaced with the average luminance level, and the background level is compensated by the difference in luminance. Note that before compression, both of these representations are error-free, that is - lossless.

The major advantage of the new scheme is now seen from Figure 8: If the upper plane is paired with LZW compression, for example, then the proposed technique

will yield significantly better compression results (as the upper plane becomes constant everywhere). On the other hand, the lower plane that is JPEG compressed will degrade some in quality as a result of compensating the luminance. However, if the two foreground colors are close in luminance, the magnitude variation in luminance in Figure 8 is much smaller than that in Figure 7, and therefore lending to much higher quality for a given compression or vice-versa.

Observe that the resulting JPEG ringing artifacts will be significantly reduced, as compared with the current approach (the Selective Model). As already explained, the current approach can only obtain similar compression results (i.e., constant upper plane) by picking a single foreground color (preferably the darkest) into the upper plane, and (a-symmetrically) alternating between the other foreground color and background in the lower plane (see Figure 7). Such schemes unnecessarily lead to larger degradation in image quality (noticeable ringing artifacts) for the same level of compression, or vice versa.

One final comment regarding the choice of foreground color is in order. Some constraints do exist, particularly when the internal bit precision is limited (to 8 bits per component, for example). In our previous example, we have selected to carry the average

foreground color in the upper plane. While this particular choice makes a lot of sense visually, it does have a drawback: some of the foreground colors will be lighter, while others will be darker.

5 As a result, the error will generally be signed. If the error becomes large in magnitude, and the background is already very light, there is risk that the result could overflow (the 8 bit range) when the error is compensated for. In our example, since the
10 error was small by construction (only 5 units in luminance (L^*), and the image was never too light to go out of range. One possible solution to this problem (without allocating more bits) is to limit the error magnitude (by reducing the foreground level) such that
15 it will be made to fit. This is another example of the greater flexibility of the proposed approach, where this time some compression is traded for keeping the representation word size fixed.

Simulation Results

20 The example described in Figure 4 was simulated using Adobe Photoshop. The pattern was created and combined with the background image "Lenna" as described above. The resulting source image (Figure 9) was used to generate the following simulation results. A few of
25 the underlying assumptions and comments are:

1. Foreground colors were entered in L*A*B*,
where "red" = (40,20,0) , "green" =
(50,-20,0).
2. The image size is 512 x 512. Listed file
5 sizes are approximate (e.g., headers,
Q-tables, markers, may be included, but the
overall effect is very small).
3. The lower plane was compressed using the
built-in JPEG, with medium-quality setting
10 (level 3).
4. The standard built-in PhotoshopTM JPEG
algorithm was used, which to the best of the
inventors' knowledge converts color data to
YCC first, and then uses a 4:2:2 subsampling
15 scheme.
5. The selector and upper planes were compressed
using the standard built-in LZW that comes
with TIFF.
6. To make the comparison fair, data that is
20 otherwise ignored was replaced by artificial
(averaged) data to yield good compression
results (see the specific explanations
below).

25 Figures 10a through 10c show the current DIR
representation, following the Selective Imaging Model.
Figure 10a shows the upper plane, composed of

alternating (noisy) stripes of red and green. Here,
data that exist where the selector plane is false
(i.e., where it is ignored) was filled with the average
color of the four surrounding 4 x 4 quadrants directly
5 above, below, and to the left/right of the quadrant to
which it belongs.

Figure 10b depicts the binary mask, where black
corresponds to "0", or false. Figure 10c describes the
lower plane. Once again, one reasonable method for
10 filling in the missing data (that has been "lifted" to
the upper plane) is to use averaging.

In this case, the original (clean) version of
"Lenna" was subsampled by a factor of 8 in each
direction, then it was expanded back to the full size
15 (using bi-cubic averaging each time). The
low-resolution data so created was then used to replace
the red/green checker pattern in the lower plane prior
to applying the JPEG compression. Note that this
method of replacing data for improving compression is
20 the cause for the noticeable edge artifacts; Whenever
the image content is relatively flat, the above
averaging operation is indistinguishable from the
background. On edges, however, this averaging tend to
blend colors together, and therefore becomes more
25 noticeably different.

Figures 11a - 11c show how the same image would be represented using the proposed Additive Imaging Model. As before, Figure 9 shows the source image for reference. Note that the selector plane (Figure 11b) is also identical to the one in Figure 10b. To demonstrate the huge advantage of the Additive Model, the number of foreground colors has been chosen to be just one color, knowing that this will compress very well in LZW. To select the one color, all the foreground red and green pixels have been averaged together to form the single dark color shown in Figure 11a. This single color was then used to fill in the complete foreground (Upper) plane.

Note that unlike the Selective Model case, no noise exists in the upper plane as a result of the above decision. Applying this decision, however, has introduced an error to the representation. This error needs to be corrected by adjusting the background (Lower) plane as previously described. For each pixel in the foreground, the error was calculated by taking the (3 dimensional) difference between the actual pixel value and the filled in foreground color, and subtracting it from the background. Note that this adjustment merely represents a transformation to shift the data between the upper and lower planes. This transformation is lossless in the sense that the

adjusted planes represent the same source image, with
no error, exactly as they did before the adjustment.
However, this transformation enables much better
compression (as the following simulation results
5 testify).

The added flexibility of the Additive Model allows
one to choose how to fine-tune the data between the
upper and lower planes as to make the overall
compression as tight as possible.

10 The resulting lower plane is given in Figure 11c.
The outcome can be reasoned as following: The
foreground red color happens to be darker than the
green (L* value of 40 versus 50). Where the red
foreground pixels used to be, the error-adjusted pixels
15 appear to be slightly darker and redder. This is due to
the fact that the average foreground color is lighter
(L* value of approximately 45), and close to neutral.
Therefore the error is making the background redder,
and slightly darker. Similarly, where the green pixels
20 used to be, the error-corrected background is made
lighter and greener. Since the image has a relatively
large red component almost everywhere, this is seen as
taking away some red (making it appear to be more
greenish-yellow, particularly as seen on the skin
25 tones).

Table 1 summarizes resulting file sizes, quality, and compression ratios obtained. For this example, the proposed Additive technique is clearly far superior (by an order of magnitude) as compared to the existing Selective approach, both in overall compression, as well as in the quality and accuracy of reconstruction.

Image/ Plane	Existing Selective Model			Proposed Additive Model		
	file size (bytes)	Comp- ression ratio	reconst- ruction quality	file size (bytes)	com- pression ratio	recon- struction quality
10 Source Image	786,432			786,432		
Upper Plane	412,601	1.91		3+header	262,144	
Selector plane	13,875	56.68		13,875	56.68	
Lower plane	49,945	15.75		53,654	14.66	
15 Total Representa- tion Size	476,421	1.651	good	67,564	11.64	excellent

Table 1

In addition to an embodiment consisting of specifically designed integrated circuits or other electronics, the present invention may be conveniently implemented using a conventional general purpose or a specialized digital computer or microprocessor programmed according to the teachings of the present disclosure, as will be apparent to those skilled in the computer art.

Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present

disclosure, as will be apparent to those skilled in the software art. The invention may also be implemented by the preparation of application specific integrated circuits or by interconnecting an appropriate network of conventional
5 component circuits, as will be readily apparent to those skilled in the art.

The present invention includes a computer program product which is a storage medium (media) having instructions stored thereon/in which can be used to program
10 a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory
15 devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

Stored on any one of the computer readable medium (media), the present invention includes software for
20 controlling both the hardware of the general purpose/specialized computer or microprocessor, and for enabling the computer or microprocessor to interact with a human user or other mechanism utilizing the results of the present invention. Such software may include, but is not
25 limited to, device drivers, operating systems, and user applications. Ultimately, such computer readable media

further includes software for performing at least one of additive model representation and reconstruction.

Included in the programming (software) of the general/specialized computer or microprocessor are software
5 modules for implementing the teachings of the present invention, including, but not limited to, separating planes of a source image, averaging at least one of foreground and background colors, replacing colors, and compensating for error introduced by color replacement in one plane by
10 feeding error into a second plane, storage, communication of results, and reconstructing an image according to the processes of the present invention.

Obviously, numerous modifications and variations of the present invention are possible in light of the above
15 teachings. It is therefore to be understood that within the scope of the appended claims, the invention may be practiced otherwise than as specifically described herein.